The Implementation of the Next Generation 64b Itanium™ Microprocessor

Samuel Naffziger
Hewlett Packard, Fort Collins, CO

Gary Hammond
Intel, Santa Clara, CA

Processor Standouts

- Leading edge processor jointly developed by two companies across multiple sites
  - This paper focuses on overview and methodologies
- Large on-chip cache contained in 4 highly optimized arrays (papers 6.6, 25.5, 6.7)
- Lowest latency cache on a mainstream CPU
  - Zero load to use penalty for 16KB L1D cache (6.6)
  - Zero bubble taken branch penalty using the same technology in L1I
- 6 general purpose, full symmetrically bypassed integer units (paper 25.6)
- 64GB/s sustainable cache bandwidth (paper 25.5)

Processor Structure

McKinley Block Diagram

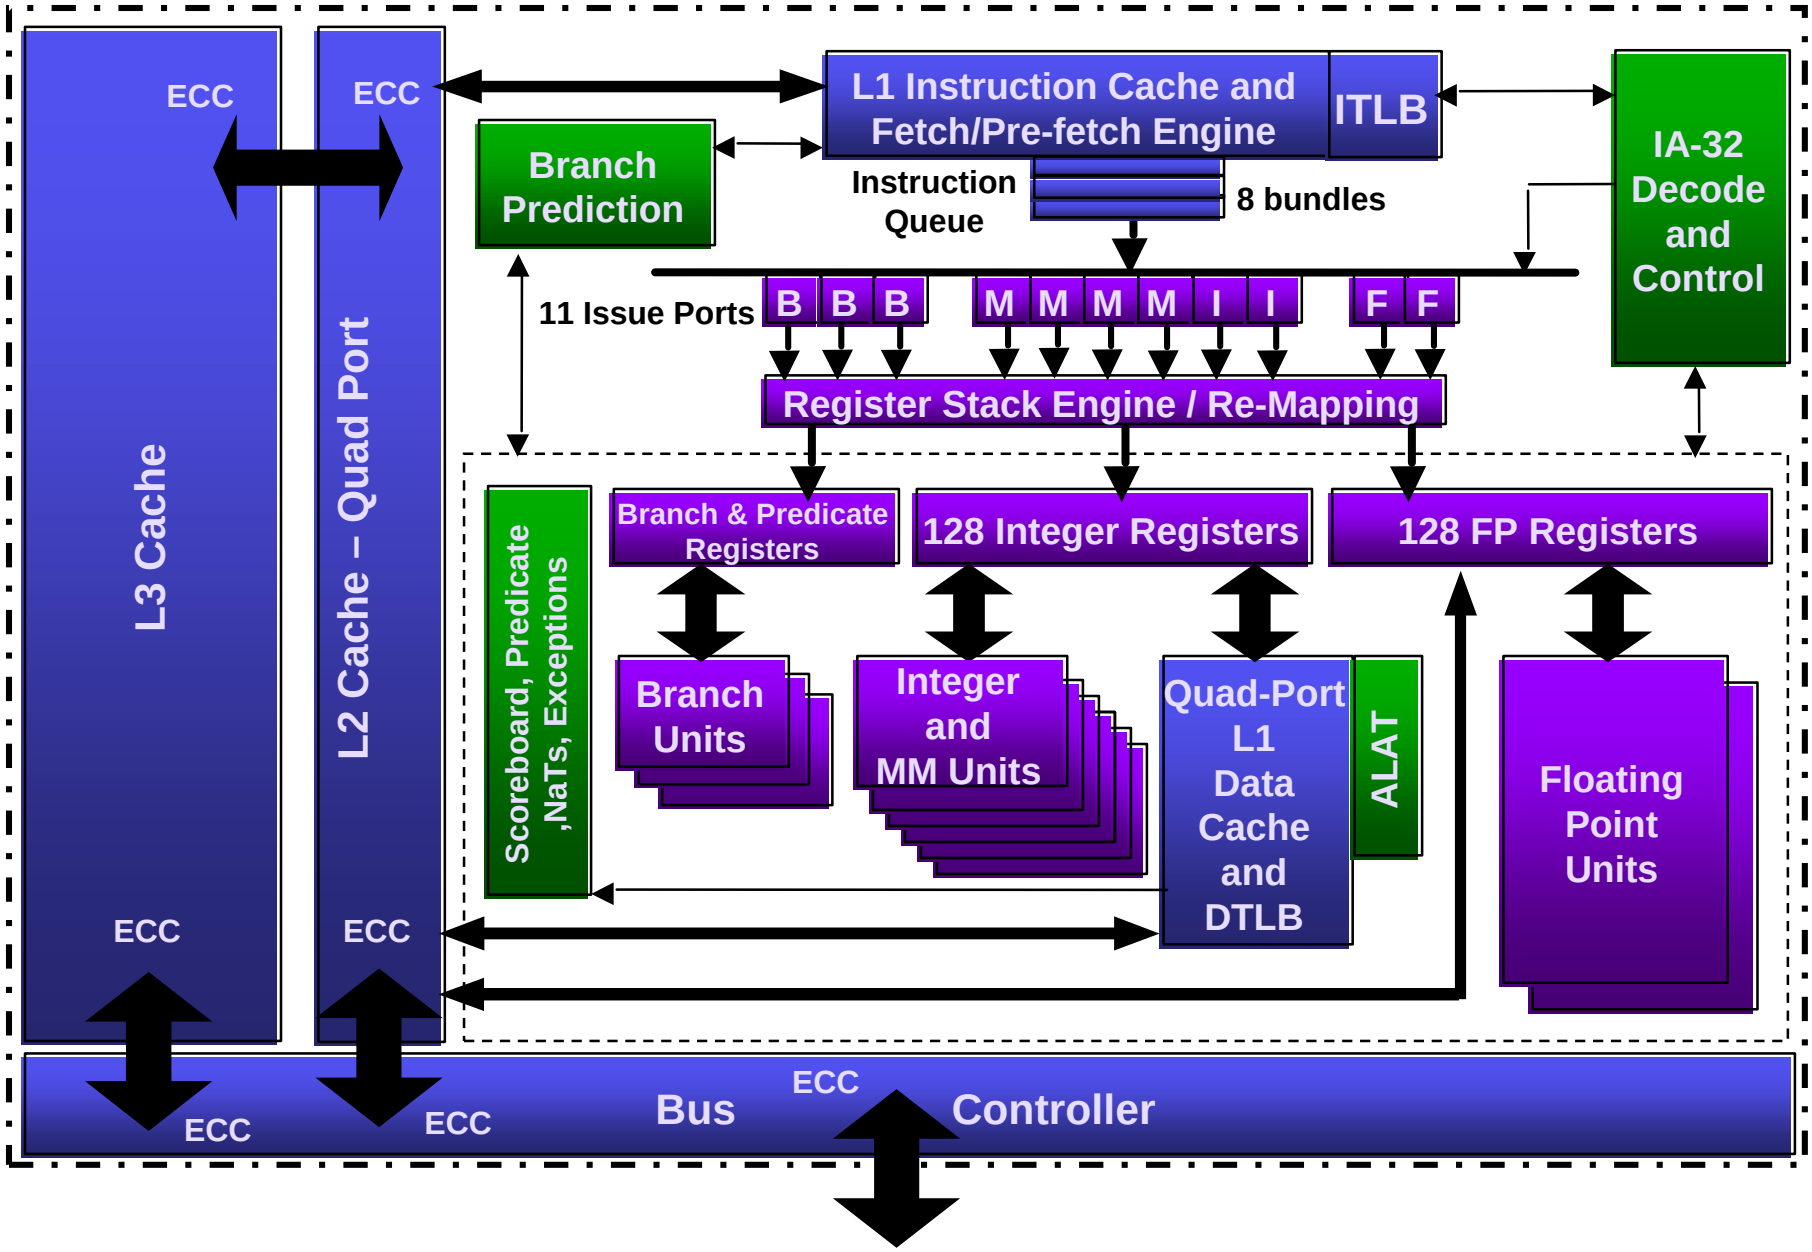

McKinley Overview

- .18µm bulk, 6 layer Al process
- 8 stage, fully stalled in-order pipeline
- Symmetric six integer-issue design
- IA32 execution engine integrated
- 3 levels of cache on-die totaling 3.3MB
- 221 Million transistors
- ECC / parity protection on all memory arrays
- 130W @1GHz, 1.5V

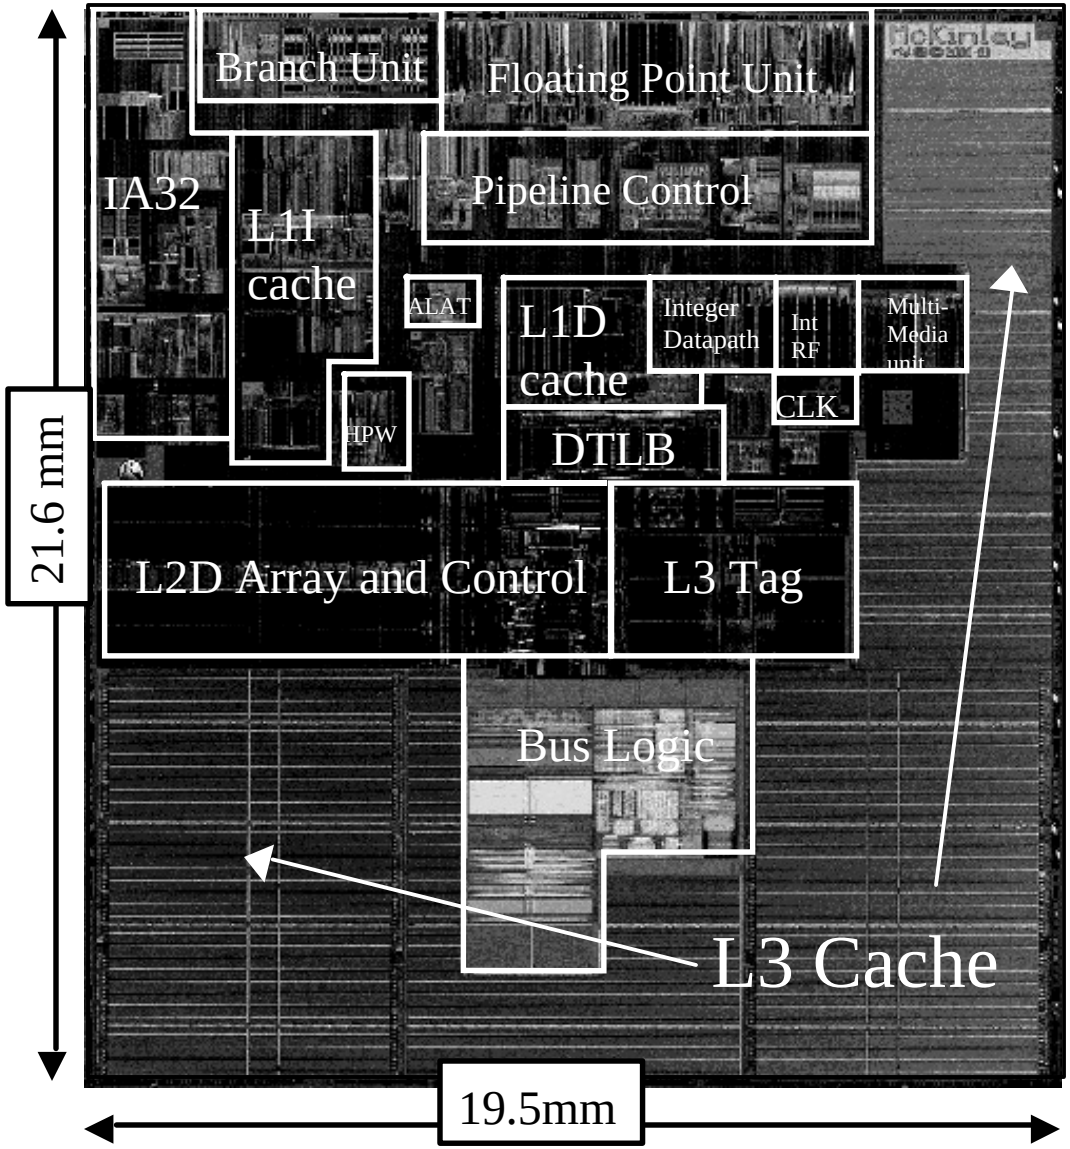

McKinley vs. other 64b .18? m processors

(sources provided in final slide)

	McKinley	USIII	EV7	Power 4 (2 cores)
Frequency (GHz)	1.0	1.05	1.0-1.2	1.3
Pipe stages (mpb)	**8 (6)**	14 (8)	9 (~11)	12 (~11)
Sustainable Int BW	**6 / cycle**	3 / cycle	4/cycle	3/cycle
FP units	2/cycle	2/cycle	2/cycle	2/cycle
On chip cache	**3.3MB** **4 arrays**	96KB 2 arrays	1.8MB 3 arrays	1.7MB 2.5 per core
D Cache read BW	**64GB/s**	16.8	19.2	41.6 / core
Die size (mm^2)	**421**	244	397	400 est.
Core size (no IO, only lowest level caches)	142	206	115	100 est.
Power (Watts)	**130**	75	125	125

Latching Methodology

- Pulse based latching (95% of static latches)
- Local and centralized pulse generation
- 85% of latches are fully scannable
- M/S version swapped in for race fixes (total of 34% done)

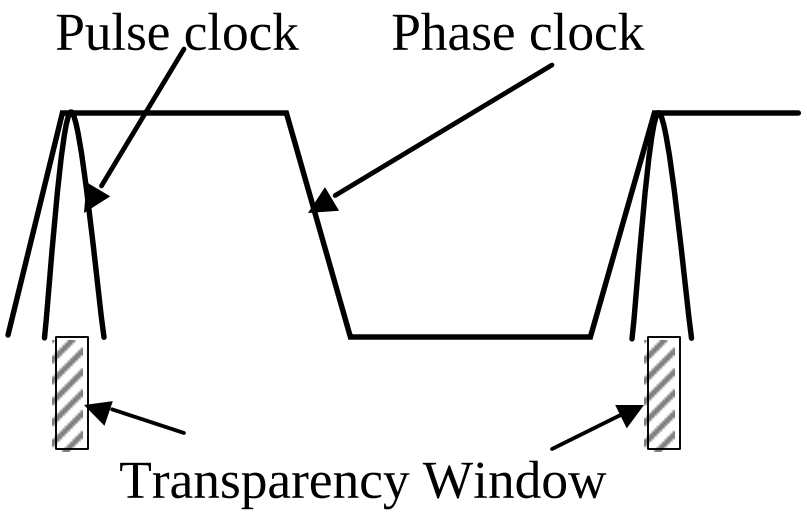

Motivations:

- **Performance**. less overhead, skew and logic imbalance compensated with transparency
- **Power**. fewer latches and clocks
- Lower **area** and **simplicity** of use

Challenges:

- **Hold time**. Back to back flops *race through* – verification and fixing failures incurs overhead
- **Pulse variation.** Requires strict control over pulses and distribution
- **Logic balancing.** Each cycle needs balancing to +/8%

Latching Methodology

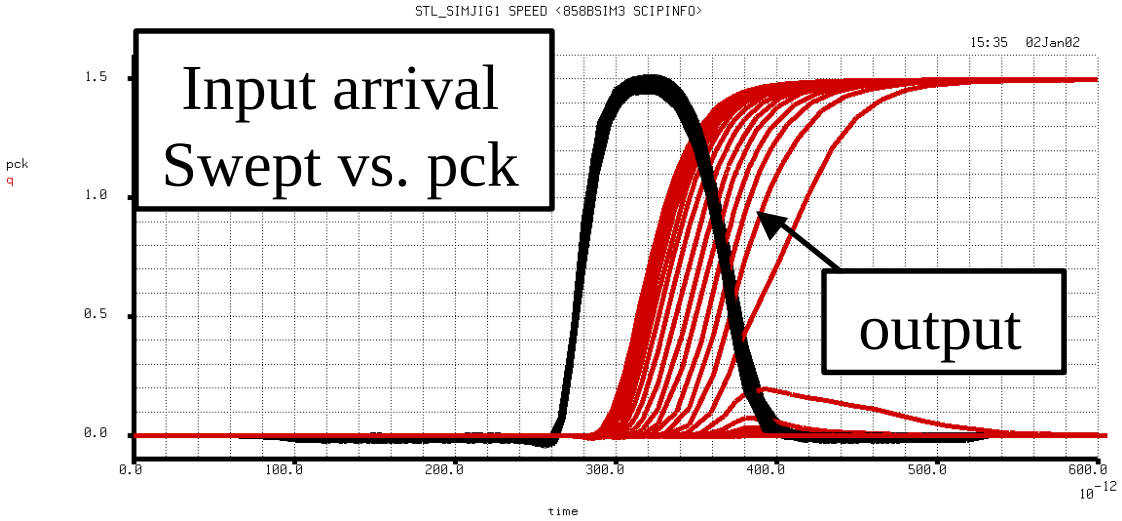

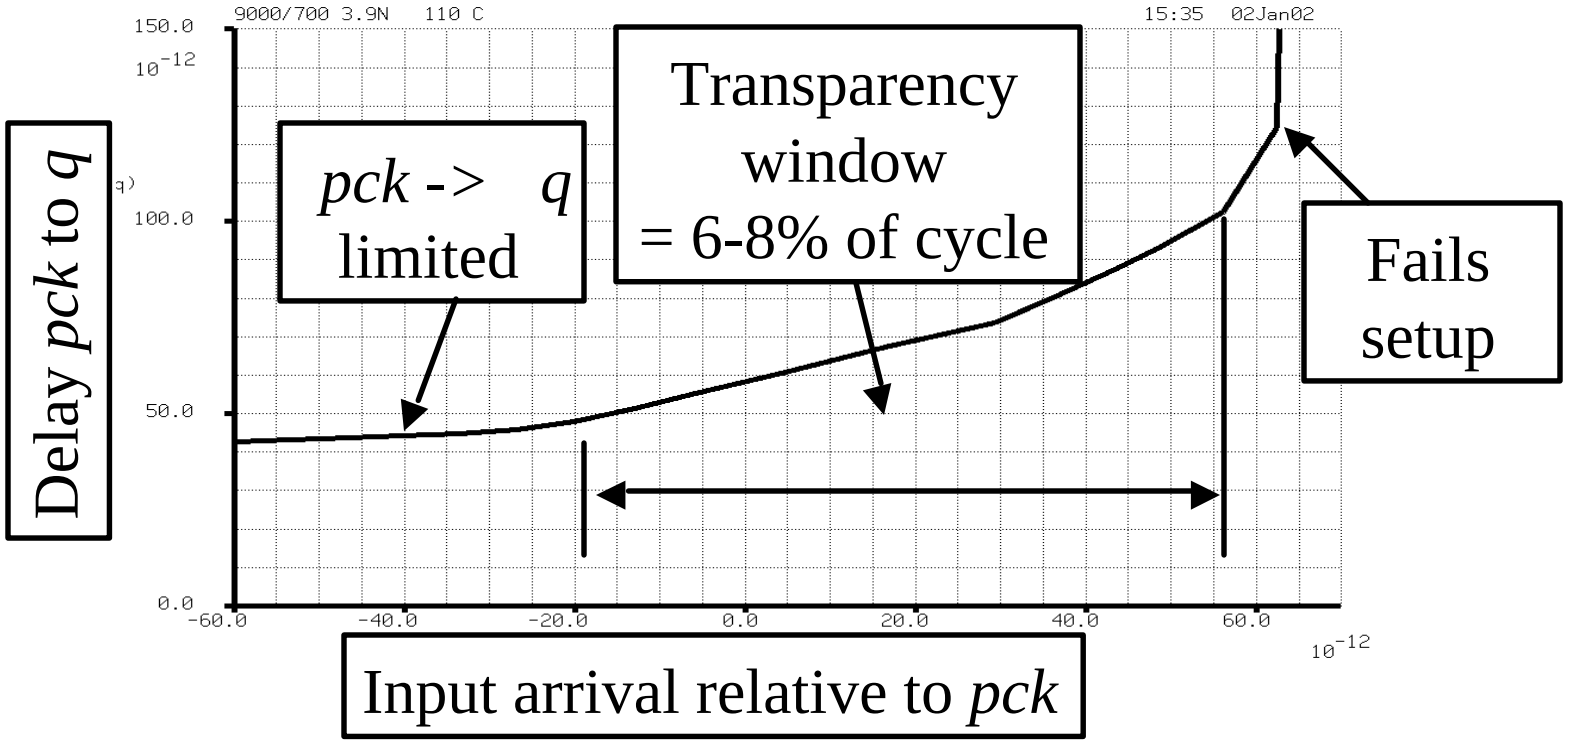

Latching Methodology

Pulse can be generated externally and distributed to multiple latches

Dynamic Scan slave for area savings

Input can be driven by any logic gate if drive strength requirements are met

Key Transistors for pulse width management

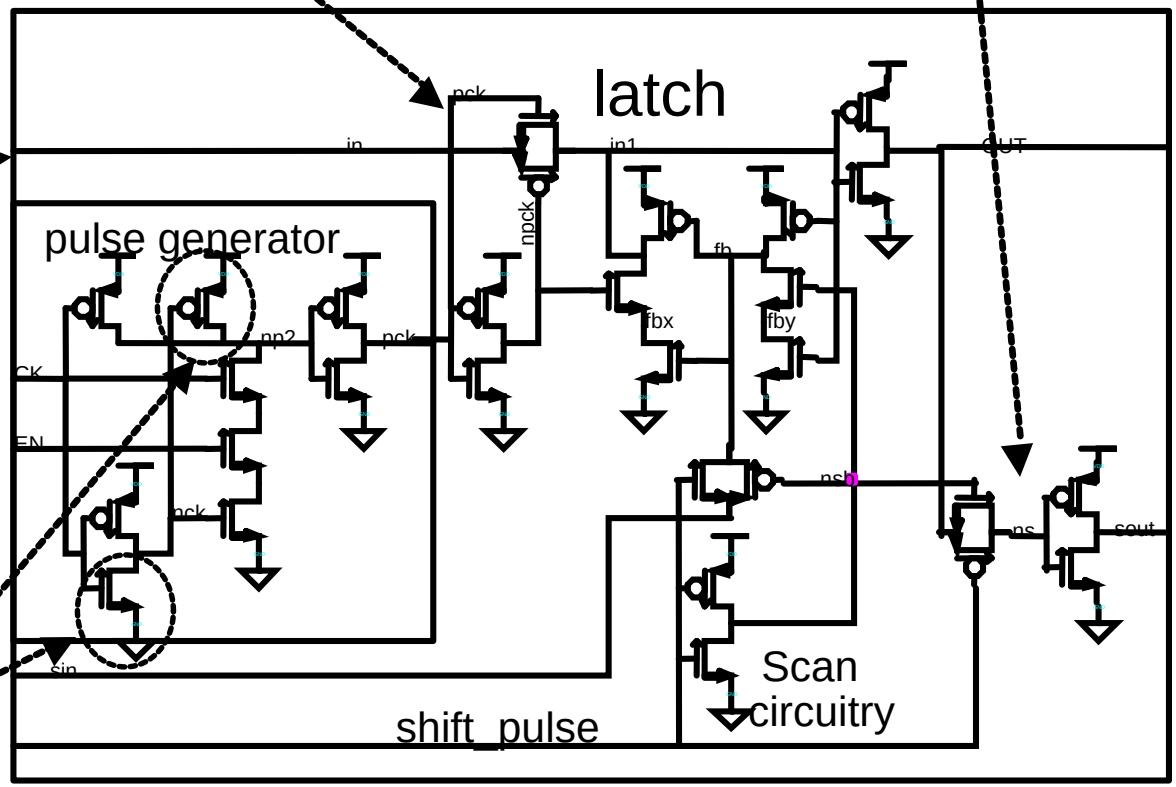

Latching Methodology Eliminating Races

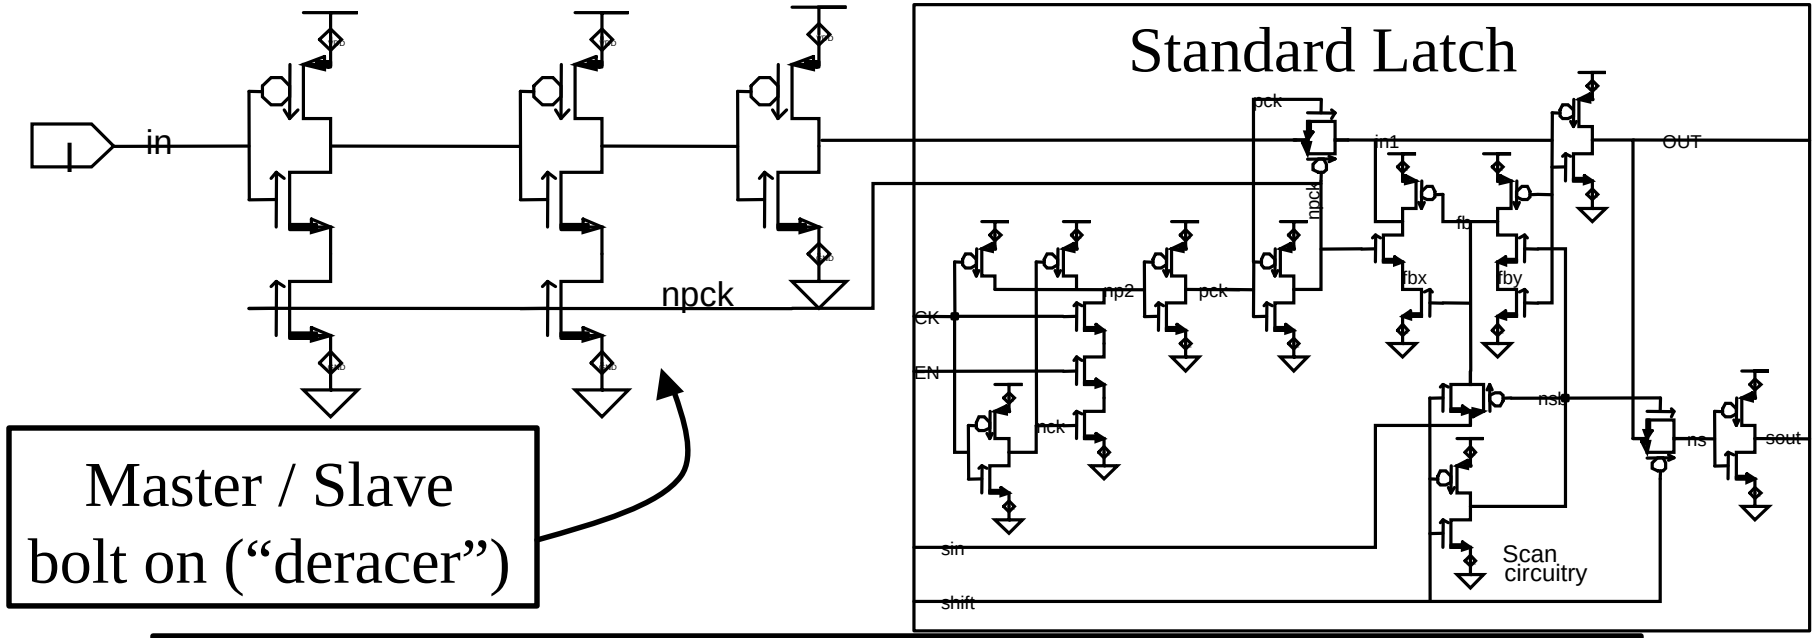

- *Low overhead* master latch. Static keepers can be removed since the clock is a pulse
- Provides *pulse width invariant* hold time benefit
- Impact to max time (flow through delay) is only 60% of the additional hold time (vs. 100% w/ a delay buffer)

Domino Latch Methodology Goals

- Integrate cleanly with static domain
- Combine high performance with ease of use
  - Phase stealing capability
  - Essentially zero latch flow through overhead
  - Integrated in with conversion of existing gates
- Enable thorough testability
- Latch monotonic hi -> low as well as lo −> hi

Domino Latch Methodology

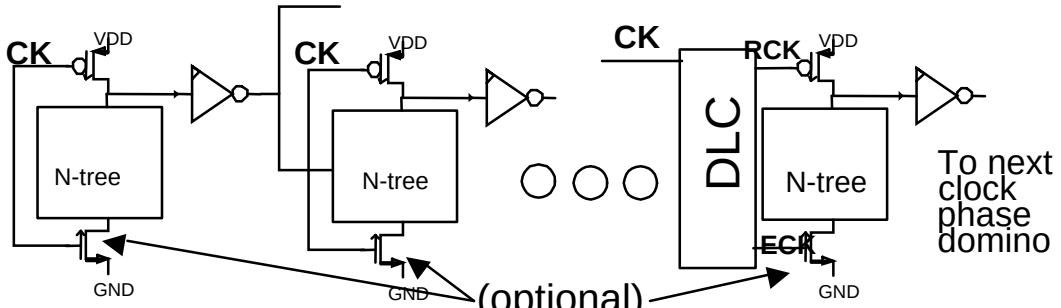

Phase clock same as used for static latches

Precharge Clock (RCK) is an inverted pulse

Evaluate Clock (ECK) is a delayed version of CK

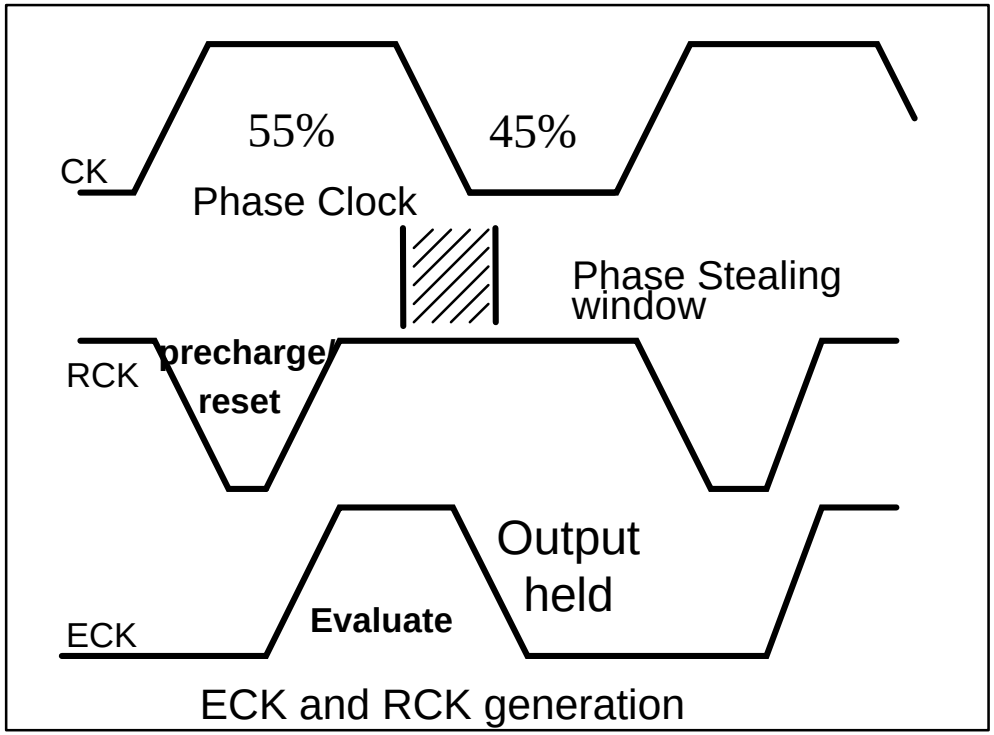

RCK & ECK derived from CK with a bolt-on circuit (DLC)

Domino Latch Methodology

DLC circuit for latching domino outputs

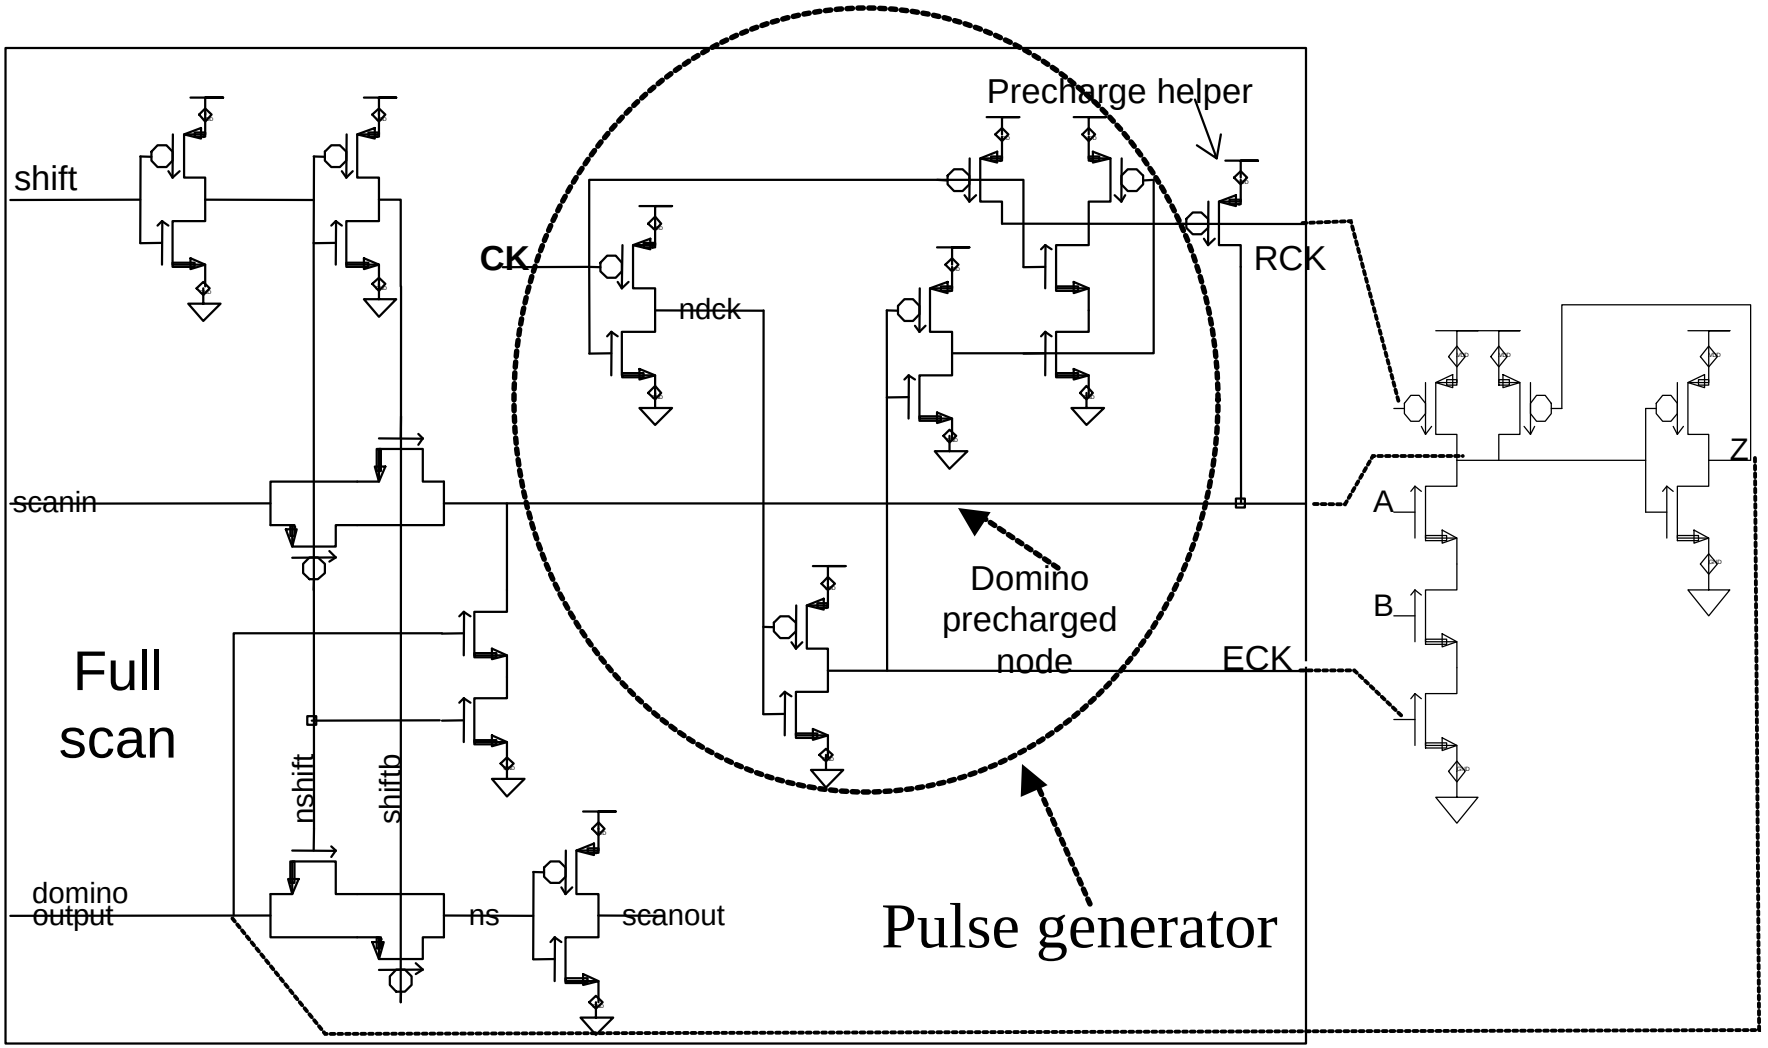

Domino Latch Methodology
Static to domino interface

- Design goals:
  - Edge triggered
  - Low setup and hold
  - Fast clock to q
  - Can incorporate logic
  - Compact

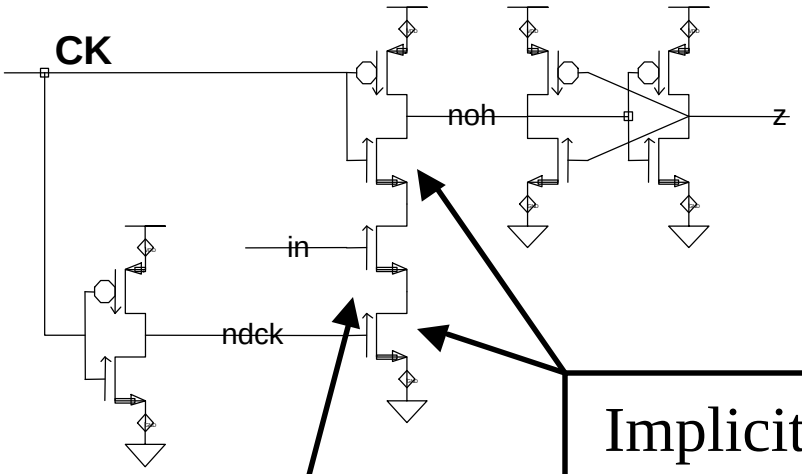

Implicit pulse

More complex pulldown tree can perform logic

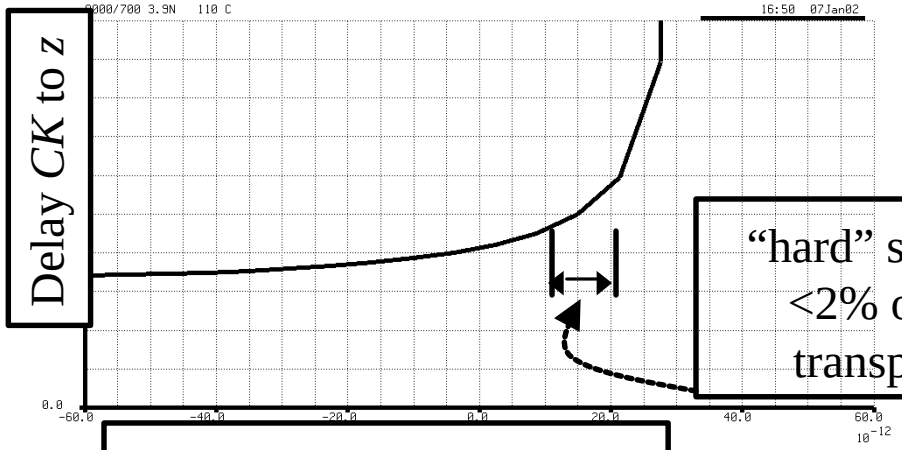

"hard" setup with <2% of cycle transparency

Input arrival relative to *CK*

Clock Management Methodology

- Hold time analysis and skews can become extremely burdensome on a design with 157K latches if not done properly. McKinley enforced several simplifications
  - Tight control of clock gaters / buffers: 24 types total and each carefully tuned in SPICE for consistency
  - Similar control of all latching elements (30 total)
  - **No** buffering of clocks by non-standard gaters allowed
  - Extracted skew of all clock nets analyzed in SPICE (custom tool) to ensure < 20ps total RC as well as load matching
- Thus clock arrival times could be idealized in timing analysis as could latch behavior
  - Ease of analysis, faster convergence, lower total skew

Latching Methodology Results

- 23544 total qualifiable clock buffers
- A total of 157,059 static latches
  - 64,391 clocked by local pulse generators
  - 52,820 are “deraced” with a M/S bolt-on or buffer
- 25,380 dynamic latches
  - 11K DLCs and 14.3K zero/one catchers
- zero hold time issues in silicon
- 2 electrical marginality issues due to insufficient drive strength into the pass gate (tool escapes)

Processor Statistics

- Over 18,200 global nets connecting up units at the top level
- ~700nF explicit bypass cap
- Power grid uniform in M5/6
  - Consumes 22% of M5 and 62% of M6
- Bus I/Os distributed on 4 stripes in the L3 cache
- 25M logic transistors, 221M total

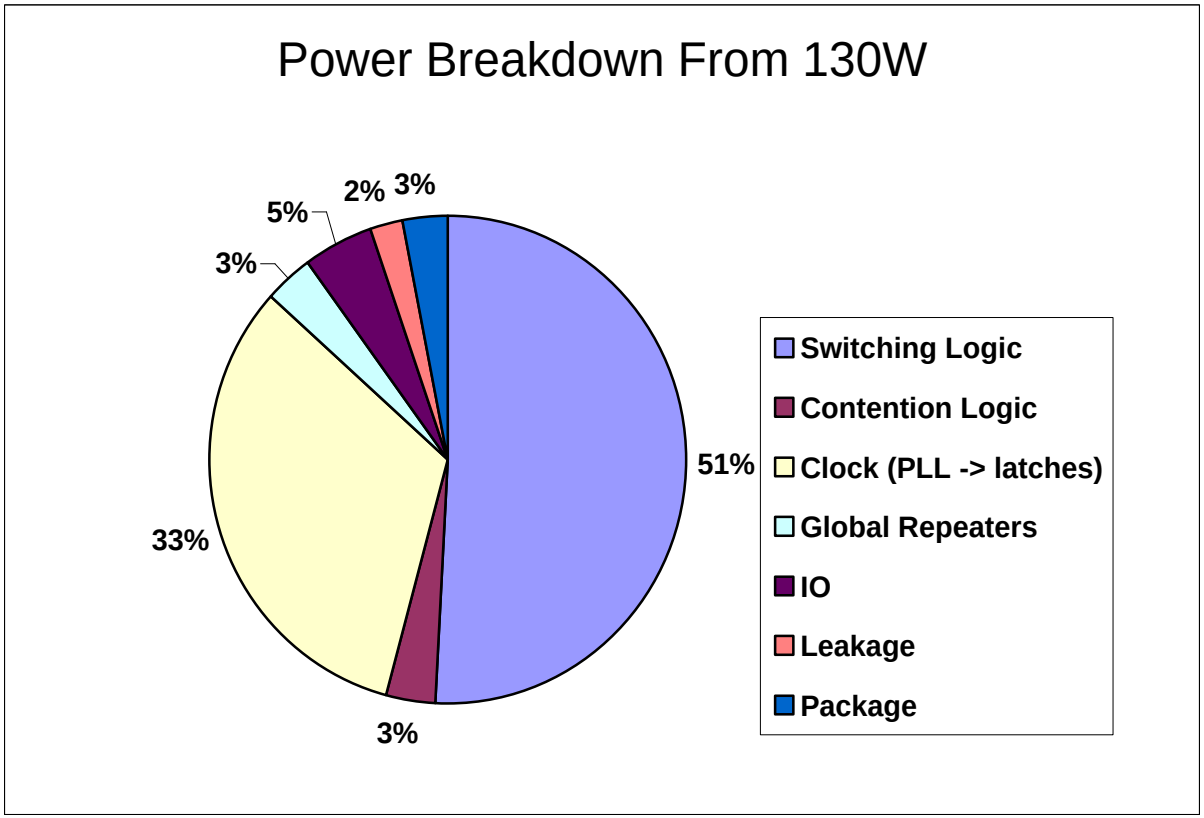

Summary

- McKinley's advanced circuit design further exploits the potential of the Itanium architecture with:
  - Large, high bandwidth integrated on-die caches (4 arrays totaling 3.3MB with 4 integer and 4 FP ports)
  - Reduced latencies (0, 5 and 12 cycle load to use penalites for L1, L2 and L3 caches respectively)
  - Increased core frequency (1 GHz)
  - Shorter pipeline
  - Greater execution bandwidth
  - Bus bandwidth (6.4GB/s system bus (128b @ 400MT/s) provides memory and IO bandwidth headroom)
- While maintaining compatibility with Itanium-based software

Competitive Data Sources

- USIII
  - ISSCC 2000, papers 25.1, 25.2
- EV7
  - ISSCC 2001, paper 15.6
  - Microprocessor Report, vol 10, Num. 14 “Digital 21264 Sets New Standard”
- Power 4
  - ISSCC 2001, paper 15.2
  - MPR vol 13, No. 13 “Power4 Focuses on Memory Bandwidth”
  - MPR 11/20/2000 “IBM’s Power4 Unveiling Continues”